arXiv:1510.04822v2 [stat.ML] 19 Oct 2015

SGD with Variance Reduction beyond Empirical Risk Minimization

Massil Achab CMAP, Ecole Polytechnique
Agathe Guilloux LSTA, Univ. Paris 6
Stéphane Gaïffas CMAP, Ecole Polytechnique
Emmanuel Bacry CMAP, Ecole Polytechnique

Abstract

We introduce a doubly stochastic proximal gradient algorithm for optimizing a finite average of smooth convex functions, whose gradients depend on numerically expensive expectations. Our main motivation is the acceleration of the optimization of the regularized Cox partial-likelihood (the core model used in survival analysis), but our algorithm can be used in different settings as well.

The proposed algorithm is doubly stochastic in the sense that gradient steps are done using stochastic gradient descent (SGD) with variance reduction, where the inner expectations are approximated by a Monte-Carlo Markov-Chain (MCMC) algorithm. We derive conditions on the MCMC number of iterations guaranteeing convergence, and obtain a linear rate of convergence under strong convexity and a sublinear rate without this assumption.

We illustrate the fact that our algorithm improves the state-of-the-art solver for regularized Cox partial-likelihood on several datasets from survival analysis.

1 Introduction

During the past decade, advances in biomedical technology have brought high dimensional data to biostatistics and survival analysis in particular. Today's challenge for survival analysis lays in the analysis of massively high dimensional (numerous covariates) and large-scale (large number of observations) data, see in particular Murdoch and Detsky (2013). Areas of application outside of biostatistics, such as economics (see Einav and Levin (2014)), or actuarial sciences (see Neumaier (2014)) are also concerned.

One of the core models of survival analysis is the Cox model (see Cox (1972)) for which we propose, in the present paper, a novel scalable optimization algorithm tuned to handle massively high dimensional and large-scale data. Survival data $(y_i, x_i, \delta_i)_{i=1}^{n_{\text{pat}}}$ contains, for each individual $i = 1, \dots, n_{\text{pat}}$, a features vector $x_i \in \mathbb{R}^d$, an observed time $y_i \in \mathbb{R}_+$, which is a failure time if $\delta_i = 1$ or a right-censoring time if $\delta_i = 0$. If $D = \{i : \delta_i = 1\}$ is the set of patients for which a failure time is observed, if $n = |D|$ is the total number of failure times, and if $R_i = \{j : y_j \geq y_i\}$ is the index of individuals still at risk at time y_i, the negative Cox partial log-likelihood writes

$$-\ell(\theta) = \frac{1}{n} \sum_{i \in D} \Big[-x_i^\top \theta + \log \Big(\sum_{j \in R_i} \exp(x_j^\top \theta) \Big) \Big] \quad (1)$$

for parameters $\theta \in \mathbb{R}^d$. With high-dimensional data, a regularization in term is added to the partial likelihood to automatically favor sparsity in the estimates, see Tibshirani (1997) and Simon et al. (2011) for a presentation of Lasso and elastic-net penalizations, see also the review paper by Witten and Tibshirani (2009) for an exhaustive presentation. Several algorithms have been proposed to solve the regularized optimization problem at hand, see Park and Hastie (2007); Sohn et al. (2009); Goeman (2010) among others. These implementations use Newton-Raphson iterations, i.e. large matrices inversions, and can therefore not handle large-scale data. Cyclical coordinate descent algorithms have since been proposed and successfully implemented in `R` packages `coxnet` and `fastcox`, see Simon et al. (2011); Yang and Zhou (2013). More recently Mittal et al. (2013) adapted the column relaxation with logistic loss algorithm of Zhang and Oles (2000) to the Cox model. The fact that all these algorithms are of cyclic coordinate descent type solve the problem, supported by Newton-Raphson type algorithms, of large matrices inversions.

Yet another computationnally costly problem, specific to the Cox model, has not been fully addressed: the presence of cumulative sums (over indices $j \in R_i$) in the Cox

Preliminary work. Under review by AISTATS 2016. Do not distribute.

partial likelihood. This problem was noticed in Mittal et al. (2013), where a numerical workaround exploiting sparsity is proposed to reduce the computational cost. The cumulative sum prevents from successfully applying stochastic gradient algorithms, which are however known for their efficiency to handle large scale generalized linear models: see for instance SAG by Schmidt et al. (2013), SAGA by Defazio et al. (2014), Prox-SVRG by Xiao and Zhang (2014) and SDCA by Shalev-Shwartz and Zhang (2012) that propose very efficient stochastic gradient algorithms with constant step-size (hence achieving linear rates).

Such recent stochastic gradient algorithms have shown that it is possible to improve upon proximal full gradient algorithms for the minimization of convex problems of the form

$$\min_{\theta\in\mathbb{R}^d} F(\theta) = f(\theta) + h(\theta) \text{ with } f(\theta) = \frac{1}{n}\sum_{i=1}^{n} f_i(\theta), \tag{2}$$

where the functions f_i are gradient-Lipschitz and h is prox-capable. These algorithms take advantage of the finite sum structure of f, by using some form of variance-reduced stochastic gradient descent. It leads to algorithms with a much smaller iteration complexity, as compared to proximal full gradient approach (FG), while preserving (or even improving) the linear convergence rate of FG in the strongly convex case. However, such algorithms are relevant when gradients ∇f_i have a numerical complexity much smaller than ∇f, such as for linear classification or regression problems, where ∇f_i depends on a single inner product $x_i^\top \theta$ between features x_i and parameters θ.

In this paper, motivated by the important example of the Cox partial likelihood (1), we consider the case where gradients ∇f_i can have a complexity comparable to the one of ∇f. More precisely, we assume that they can be expressed as expectations, under a probability measure π_θ^i, of random variables $G_i(\theta)$, i.e.,

$$\nabla f_i(\theta) = \mathbb{E}^{G_i(\theta)\sim\pi_\theta^i}[G_i(\theta)]. \tag{3}$$

This paper proposes a new doubly stochastic proximal gradient descent algorithm (2SVRG), that leads to a low iteration complexity, while preserving linear convergence under suitable conditions for problems of the form (2) + (3).

Our main motivation for considering this problem is to accelerate the training-time of the the penalized Cox partial-likelihood, see Cox (1972) and Tibshirani (1997). The function $-\ell(\theta)$ is convex (as a sum of linear and log-sum-exp functions, see Chapter 3 of Boyd and Vandenberghe (2004), and fits in the setting (2) + (3). Indeed, fix $i \in D$ and introduce

$$f_i(\theta) = -x_i^\top\theta + \log\Big(\sum_{j\in R_i}\exp(x_j^\top\theta)\Big),$$

so that

$$\nabla f_i(\theta) = -x_i + \sum_{j\in R_i} x_j \pi_\theta^i(j)$$

where

$$\pi_\theta^i(j) = \frac{\exp(x_j^\top\theta)}{\sum_{j'\in R_i}\exp(x_{j'}^\top\theta)}, \quad \forall j \in R_i.$$

This entails that $\nabla f_i(\theta)$ satisfies (3) with $G_i(\theta)$ a random variable valued in $\{-x_i + x_j : j \in R_i\}$ and such that

$$\mathbb{P}(G_i(\theta) = -x_i + x_j) = \pi_\theta^i(j)$$

for $j \in R_i$. Note that the numerical complexity of ∇f_i can be comparable to the one of ∇f, when y_i is close to $\min_i y_i$ (recalling that $R_i = \{j : y_j \geq y_i\}$). Note also that a computational trick allows to compute $\nabla f(\theta)$ with a complexity $O(nd)$, see Appendix B for details. This makes this setting quite different from the usual case of empirical risk minimization (linear regression, logistic regression, etc.), where all the gradients ∇f_i share the same low numerical cost.

2 Comparison with previous work

SGD techniques. Recent proximal stochastic gradient descent algorithms by Defazio et al. (2014), Xiao and Zhang (2014), Shalev-Shwartz and Zhang (2012) and Schmidt et al. (2013) build on the idea of Robbins and Monro (1951) and Kiefer and Wolfowitz (1952). Such algorithms are designed to tackle large-scale optimization problems (n is large), where it is assumed implicitly that the ∇f_i (smooth gradients) have a low computational cost compared to ∇f, and where h is eventually non-differentiable and is dealt with using a backward or projection step using its proximal operator.

The principle of SGD is, at each iteration t, to sample uniformly at random an index $i \sim \mathcal{U}[n]$, and to apply an update step of the form

$$\theta^{t+1} \leftarrow \theta^t - \gamma_t \nabla f_i(\theta^t).$$

This step is based on an unbiased but very noisy estimate of the full gradient ∇f, so the choice of the step size γ_t is crucial since it has to be decaying to curb the variance introduced by random sampling (excepted for averaged SGD in some particular cases, see Bach and Moulines (2013)). This tends to slow down convergence to a minimum $\theta^\star \in \operatorname{argmin}_{\theta\in\mathbb{R}^d} f(\theta)$. Gradually reducing the variance of ∇f_i for $i \sim \mathcal{U}[n]$ as an approximation of

∇f allows to use larger – even constant – step sizes and to obtain faster convergence rates. This is the underlying idea of two methods introduced recently in Defazio et al. (2014), Xiao and Zhang (2014), that use updates of the form

$$w^{t+1} \leftarrow \theta^t - \gamma\Big(\nabla f_i(\theta^t) - \nabla f_i(\tilde{\theta}) + \frac{1}{n}\sum_{j=1}^{n} \nabla f_j(\tilde{\theta})\Big),$$

and $\theta^{t+1} \leftarrow \text{prox}_{\gamma h}(w^{t+1})$. In Xiao and Zhang (2014), $\tilde{\theta}$ is fully updated after a certain number of iterations, called *phases*, whereas in Defazio et al. (2014), $\tilde{\theta}$ is partially updated after each iteration. Both methods use stochastic gradient descent steps, with variance reduction obtained via the centered control variable $-\nabla f_i(\tilde{\theta}) + \frac{1}{n}\sum_{j=1}^{n} \nabla f_j(\tilde{\theta})$, and achieve linear convergence when F is strongly-convex, namely $\mathbb{E}F(\theta^k) - \min_{\in \mathbb{R}^d} F(x) = O(\rho^k)$ with $\rho < 1$, which make these algorithms state-of-the-art for many convex optimization problems.

Numerically hard gradients. A very different, nevertheless classical, "trick" to reduce the complexity of the gradient computation, is to express it, whenever the statistical problem allows it, as the expectation, with respect to a non-uniform distribution π_θ, of a random variable $G(\theta)$, i.e., $\nabla f(\theta) = \mathbb{E}^{G(\theta)\sim\pi_\theta}[G(\theta)]$. Optimization problems with such a gradient have generated an extensive literature from the first works by Robbins and Monro (1951), and Kiefer and Wolfowitz (1952). Some algorithms are designed to construct stochastic approximations of the sub-gradient of $f + h$, see Nemirovski et al. (2009); Juditsky and Nemirovski (2010); Lan (2010); Duchi et al. (2011). Others are based on proximal operators to better exploit the smoothness of f and the properties of h, see Hu et al. (2009); Xiao (2009); Atchade et al. (2014). In this paper, we shall focus on the second kind of algorithms. Indeed, our approach is closer to the one developed in Atchade et al. (2014), though, as opposed to ours, the algorithm developed in this latter work is based on proximal full gradient algorithms (not doubly stochastic as ours) and does not guarantee a linear convergence.

Our setting. The setting of our paper is original in the sense that it combines both previous settings. As in the stochastic gradient setting, the gradient can be expressed as the sum of n components, where n can be very large. However, since these components are time-consuming to compute directly, following the expectation based gradient computation setting, they are expressed as averaged values of some random variables. More precisely, the gradient $\nabla f_i(\theta)$ is replaced by an approximation $\widehat{\nabla} f_i(\theta)$ obtained by an MCMC algorithm. Our algorithm is, to the best of our knowledge, the first one to propose a combination of two stochastic approximations in this way, hence the name *doubly stochastic*, which allow to deal with both, eventual large values for n and the inner complexity of each gradient ∇f_i computation.

The idea to mix SGD and MCMC has also been raised recently in the very different setting of *implicit* stochastic gradient descent, see Toulis and Airoldi (2014). Note also that in our approach we make two stochastic approximations to the gradient using random training points, while the doubly stochastic approach from Dai et al. (2014) performs two stochastic approximations to the gradient using random training points and random features for kernel methods.

3 A doubly stochastic proximal gradient descent algorithm

Following the ideas presented in the previous section, we design a *doubly stochastic proximal gradient descent algorithm* (2SVRG), by combining a variance reduction technique for SGD given by Prox-SVRG Xiao and Zhang (2014), and a Monte-Carlo Markov-Chain algorithm to obtain an approximation of the gradient $\nabla f_j(\theta)$ at each step. Thus, in the considered setting the full gradient writes

$$\nabla f(\theta) = \mathbb{E}^{i\sim\mathcal{U}}[\nabla f_i(\theta)] = \mathbb{E}^{i\sim\mathcal{U}}\, \mathbb{E}^{G_i(\theta)\sim\pi_\theta^i}[G_i(\theta)],$$

where $\mathcal{U}$ is the uniform distribution on $\{1, \dots, n\}$, so our algorithm contains two levels of stochastic approximation: uniform sampling of i (the variance-reduced SGD part) for the first expectation, and an approximation of the second expectation w.r.t π_θ^i by means of Monte-Carlo simulation. The 2SVRG algorithm is described in Algorithm 1.

Following Prox-SVRG by Xiao and Zhang (2014), this algorithm decomposes in *phases*: iterations within a phase apply variance reduced stochastic gradient steps (with a backward proximal step, see lines 7 and 8 in Algorithm 1). At the end of a phase, a full-gradient is computed (lines 10, 11) and used in the next phase for variance reduction. In practice, the datasets are finite and a full-gradient is computable, even if it is said to be *numerically hard*. Within a phase, each inner iteration samples uniformly at random an index i (line 4) and obtains an approximation of the gradient ∇f_i at the previous iterate θ^t by applying N_k iterations of a Monte-Carlo Markov-Chain (MCMC) algorithm.

Intuitively, the sequence N_k should be increasing with the phase number k, as we need more and more precision as the iterations goes on (this shall be confirmed in Section 4). The important point of our algorithm resides

Algorithm 1 Doubly stochastic proximal gradient descent (2SVRG)

Require: Number of phases $K \geq 1$, phase-length $m \geq 1$, step-size $\gamma > 0$, MCMC number of iterations per phase $(N_k)_{k=1}^K$, starting point $\theta^0 \in \mathbb{R}^d$

1: **Initialize:** $\tilde{\theta} \leftarrow \theta^0$ and compute $\nabla f_i(\tilde{\theta})$ for $i = 1, \dots, n$
2: **for** $k = 1$ **to** K **do**
3: **for** $t = 0$ **to** $m - 1$ **do**
4: Pick $i \sim \mathcal{U}[n]$
5: $\widehat{\nabla} f_i(\theta^t) \leftarrow$ approx. $\nabla f_i(\theta^t)$ using N_k MCMC iter.
6: $d^t = \widehat{\nabla} f_i(\theta^t) - \nabla f_i(\tilde{\theta}) + \frac{1}{n}\sum_{j=1}^n \nabla f_j(\tilde{\theta})$
7: $\omega^{t+1} \leftarrow \theta^t - \gamma d^t$
8: $\theta^{t+1} \leftarrow \mathrm{prox}_{\gamma h}(\omega^{t+1})$
9: **end for**
10: Update $\tilde{\theta} \leftarrow \frac{1}{m}\sum_{t=1}^m \theta^t$, $\theta^0 \leftarrow \tilde{\theta}$, $\tilde{\theta}^k \leftarrow \tilde{\theta}$
11: Compute $\nabla f_i(\tilde{\theta})$ for $i = 1, \dots, n$
12: **end for**
13: **Return:** $\tilde{\theta}^K$

precisely in this aspect: very noisy estimates can be used in the early phases of the algorithm, hence allowing for an overall low complexity as compared to a full gradient approach.

When the π_θ^i are Gibbs probability measures, as for the previously described Cox partial log-likelihood (but for other models as well, such as Conditional Random Fields, see Lafferty et al. (2001)), one can apply Independent Metropolis-Hastings (IMH), see Algorithm 2 below, to obtain approximations $\widehat{\nabla} f_i$ of the gradients. In this case the produced chain is geometrically uniformly ergodic, see Robert and Casella (2004), and therefore meets the general assumptions required in our results (see Proposition 4.3 below). The IMH algorithm uses a proposal distribution Q which is independent of the current state u^l of the Markov chain. In the case of the Cox partial log-likelihood, at iteration t of phase k of Algorithm 1, we set $\pi = \pi_{\theta^t}^i$, and Q to be the uniform distribution over the set R_i.

Algorithm 2 Independent Metropolis-Hastings (IMH)

Require: Proposal distribution Q, starting point $u^0 \in \mathbb{R}^d$, stationary distribution π

for $l = 1, \dots, N_k$ **do**

1. **Generate:** $v^{l+1} \sim Q(\cdot)$.
2. **Update:** $\alpha = \min\left(\frac{\pi(v^{l+1})Q(u^l)}{\pi(u^l)Q(v^{l+1})}, 1\right)$.
3. **Take:** $u^{l+1} = \begin{cases} v^{l+1} & \text{with probability } \alpha \\ u^l & \text{otherwise.} \end{cases}$

end for

Return: $\bar{u} = \frac{1}{N_k}\sum_{l=1}^{N_k} u^l$

Section 4 below gives theoretical guarantees for Algorithm 1: linear convergence under strong-convexity of F is given in Theorem 4.1, and a convergence without strong convexity is given in Theorem 4.2. This improves the proximal stochastic gradient method of Atchade et al. (2014), where the best case rate is $O(1/k^2)$ using Fista (see Beck and Teboulle (2009)) acceleration scheme. Numerical illustrations are given in Section 5, where a fair comparison between several state-of-the-art algorithms is proposed.

4 Theoretical guarantees

Definitions. All the functions f_i and h are proper convex lower-semicontinuous on $\mathbb{R}^d$. The norm $\|\cdot\|$ stands for the Euclidean norm on $\mathbb{R}^d$. A function $f : \mathbb{R}^d \to \mathbb{R}$ is L-smooth if it is differentiable and if its gradient is L-Lipschitz, namely if $\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|$ for all $x, y \in \mathbb{R}^d$. A function $f : \mathbb{R}^d \to \mathbb{R}$ is μ-strongly convex if $f(x + y) \geq f(x) + \nabla f(x)^\top y + \frac{\mu}{2}\|y\|^2$ for all $x, y \in \mathbb{R}^d$ i.e. if $f - \frac{\mu}{2}\|\cdot\|^2$ is convex. The proximal operator of $h : \mathbb{R}^d \to \mathbb{R}$ is uniquely defined by $\mathrm{prox}_h(x) = \mathrm{argmin}_{y \in \mathbb{R}^d}\{h(y) + \frac{1}{2}\|x - y\|^2\}$.

Notations. We denote by i_t the index randomly picked at the t^{th} iteration, see line 4 in Algorithm 1. We introduce the error of the MCMC approximation $\eta^t = \widehat{\nabla} f_{i_t}(\theta^{t-1}) - \nabla f_{i_t}(\theta^{t-1})$ and the filtration $\mathcal{F}_t = \sigma(\theta^0, i_1, \theta^1, \dots, i_t, \theta^t)$. In order to analyze the descent steps, we need different expectations: $\mathbb{E}_t$ the expectation w.r.t the distribution of the pair $(i_t, \widehat{\nabla} f_{i_t}(\theta^{t-1}))$ conditioned on $\mathcal{F}_{t-1}$, and $\mathbb{E}$ the expectation w.r.t all the random iterates (i_t, θ^t) of the algorithm. We also denote $\theta^* = \mathrm{argmin}_{\theta \in \mathbb{R}^d} F(\theta)$.

Assumptions.

Assumption 1. We consider $F = f + h$ where $f = \frac{1}{n}\sum_{i=1}^n f_i$, with each f_i being convex and L_i-smooth, $L_i > 0$, and h a lower semi-continuous and closed convex function. We denote $L = \max_{1 \leq i \leq n} L_i$. We assume that there exists $B > 0$ such that the iterates θ^t satisfy $\sup_{t \geq 0} \|\theta^t - \theta^*\| \leq B$.

Assumption 2. We assume that the bias and the expected squared error of the Monte Carlo estimation can be bounded in the following way:

$$\|\mathbb{E}_t \eta^t\| \leq \frac{C_1}{N_k} \text{ and } \mathbb{E}_t\|\eta^t\|^2 \leq \frac{C_2}{N_k} \tag{4}$$

for the iterations t belonging to the k-th phase, where N_k is the number of iterations of the Markov chain used for the computation of $\widehat{\nabla} f_{i_t}(\theta^t)$ during phase k (see line 5 of Algorithm 1), and where C_1 and C_2 are positive constants.

Let us point out that Proposition 4.3 below gives a sufficient condition for Assumption 2 to hold.

Theorems. The theorems below provide upper bounds on the distance to the minimum in the strongly convex case, see Theorem 4.1 and in the convex case, see Theorem 4.2.

Theorem 4.1. *Suppose that $F = f + h$ is μ-strongly convex. Consider Algorithm 1, with a phase length m and a step-size $\gamma \in (0, \frac{1}{16L})$ satisfying*

$$\rho = \frac{1}{m\gamma\mu(1-8L\gamma)} + \frac{8L\gamma(1+1/m)}{1-8L\gamma} < 1. \qquad (5)$$

Then, under Assumption 1 and Assumption 2, we have:

$$\mathbb{E}[F(\tilde{\theta}^K)] - F(\theta^*) \leq \rho^K \Big(F(\theta^0) - F(\theta^*) + \sum_{l=1}^{K} \frac{D}{\rho^l N_l} \Big), \qquad (6)$$

where $D = \frac{3\gamma C_2 + BC_1}{1-8L\gamma}$.

In Theorem 4.1, the choice $N_k = k^\alpha \rho^{-k}$ with $\alpha > 1$ gives

$$\mathbb{E}[F(\tilde{\theta}^K)] - F(\theta^*) \leq D' \rho^K$$

where $D' = F(\theta^0) - F(\theta^*) + D \sum_{k \geq 1} k^{-\alpha}$ and $D > 0$ is a numerical constant. This entails that 2SVRG achieves a *linear rate* under strong convexity.

Remark 1 (An important remark). The number N_k of MCMC iterations is growing quickly with the phase number k. So, we use in practice an hybrid version of 2SVRG called *HSVRG*: 2SVRG is used for the first phases (usally 4 or 5 phases in our experiments), and as soon as N_k exceeds n, we switch to a mini-batch version of Prox-SVRG (SVRG-MB), see Nitanda (2014). A precise description of HSVRG is given in Algorithm 3 from Section 5 below. Note that overall linear convergence of HSVRG is still guaranteed, since both 2SVRG and SVRG-MB decrease linearly the objective from one phase to the other.

Theorem 4.2. *Consider Algorithm 1, with a phase length m and a step-size $\gamma \in (0, \frac{1}{8L(2m+1)})$. Then, under Assumption 1 and Assumption 2, we have*:

$$\mathbb{E}[F(\bar{\theta}^K)] - F(\theta^*) \leq \frac{D_1}{K} + \frac{D_2}{K} \sum_{k=1}^{K+1} \frac{1}{N_k}, \qquad (7)$$

where D_1 and D_2 depend on the constants of the problem, and where $\bar{\theta}^K$ is the average of iterates $\tilde{\theta}^k$ until phase K.

In Theorem 4.2, the choice $N_k = k^\alpha$ with $\alpha > 1$ gives

$$\mathbb{E}[F(\bar{\theta}^K)] - F(\theta^*) \leq \frac{D_3}{K}$$

for a constant $D_3 > 0$. This result is an improvement of the Stochastic Proximal Gradient algorithm from Atchade et al. (2014) since it is not necessary to design a weighted averaged but just a simple average to reach the same convergence rate. Also, it provides a convergence guarantee for the non-strongly convex case, which is not proposed in Xiao and Zhang (2014).

Theorems 4.1 and 4.2 show a trade-off between the linear convergence of the variance-reduced stochastic gradient algorithm and the MCMC approximation error. The next proposition proves that Algorithm 2 satisfies Assumption 2 under a general assumption on the proposal and the stationary distribution.

Proposition 4.3. *Suppose that there exists $M > 0$ such that the proposal Q and the stationary distribution π satisfy $\pi(x) \leq MQ(x)$, for all x in the support of π. Then, the error η^t obtained by Algorithm 2 satisfies Assumption 2.*

Remark 2 (Specifics for the Cox partial likelihood). Note that the assumptions required in Proposition 4.3 are met for the Cox partial likelihood: in this case, a simple choice is $M = n \max_{x \in \text{supp}(\pi)} \pi(x)$, and the Monte Carlo error η^t induced by computing the gradient of f_i at phase k using Algorithm 2 satisfies (4) with

$$C_1 = \frac{2}{|R_i|} \max_{j \in R_i} \pi^i_{\theta^{t-1}}(j)$$
$$C_2 = 36 \mathcal{C}_2 C_1^2 (1 + C_1) \max_{j \in R_i} \|x_j\|_2^2,$$

where $\mathcal{C}_2$ is the Rosenthal constant of order 2, see Proposition 12 in Fort and Moulines (2003).

5 Numerical experiments

We compare several solvers for the minimization of the objective given by an elastic-net penalization of the Cox partial likelihood

$$F(\theta) = -\ell(\theta) + \lambda \Big(\alpha \|\theta\|_1 + \frac{1-\alpha}{2} \|\theta\|_2^2 \Big),$$

where we recall that the partial likelihood ℓ is defined in Equation (1) and where $\lambda > 0$ and $\alpha \in [0, 1]$ are tuning parameters.

A fair comparison of algorithms. The doubly stochastic nature of the considered algorithms makes it hard to compare them to batch algorithms in terms of iteration number or epoch number (number of full passes over the data), as this is usually done for SGD-based algorithm. Hence, we proceed by plotting the evolution of $F(\tilde{\theta}) - F(\theta^*)$ (where $\theta^* \in \text{argmin}_{u \in \mathbb{R}^d} F(u)$ and $\tilde{\theta}$ is the current iterate of a solver) as a function of the number of inner products between a feature vector x_i and θ, effectively computed by each algorithm, to obtain the current $\tilde{\theta}$. This gives a fair way of comparing the effective complexity of all algorithms.

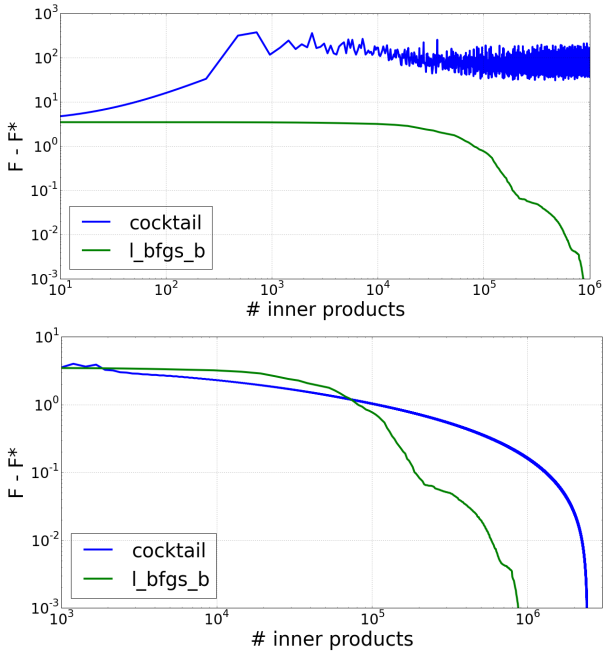

Figure 1: Convergence of Cocktail and L-BFGS-B on Lymphoma dataset. *Top:* the starting point is $\theta^0 = \mathbf{0} \in \mathbb{R}^d$. *Bottom:* the starting point is $\theta^0 = \widehat{\theta}^{(l)}$ (solution to the same objective with a slightly larger λ). This illustrates the fact that Cocktail cannot minimize directly a single objective (with a fixed λ) and requires to compute the full path of solution to converge.

About the baselines specific to the Cox model. State-of-the-art algorithms to fit the elastic-net penalized Cox partial likelihood are `cocktail` by Yang and Zhou (2013) and `coxnet`, by Simon et al. (2011). Both algorithms are combining the ideas of coordinate descent and majoration-minimization. Full convergence results for these algorithms have not yet been established, although `Cocktail` has a coordinate-wise descent property.

These algorithms however need a good starting point (near the actual minimizer) to achieve convergence (this fact is due to a diagonal approximation of the Hessian matrix, see Tibshirani and Hastie (1990), Chapter 8.). They are therefore tuned to provide good path of solutions while varying by small steps the penalization parameter λ. Indeed in this case, this starting point is naturally set at the minimizer at the previous value of λ, when minimizing along a path but cannot be guessed outside of a path. We illustrate this fact on Figure 1, where the convergence of `Cocktail` and L-BFGS-B algorithms are compared for two starting points θ_0. Experiences (not shown) showed the same behavior for `coxnet` algorithm.

Even when the starting point is set to the previous minimizer (second case in Figure 1, `cocktail`'s convergence is slower than the one of L-BFGS-B. As a consequence, we decided that no fair comparison could be conducted with `cocktail` and `coxnet` algorithms.

Baselines. We describe in this paragraph the algorithm that we put in competition in our experiments.

FISTA. This is accelerated proximal gradient from Beck and Teboulle (2009) with backtracking linesearch. Inner products necessary inside the backtracking are counted as well.

L-BFGS-B. A state-of-the-art quasi-Newton solver which provides a usually strong baseline for many batch optimization algorithms, see Liu and Nocedal (1989). We use the original implementation of the algorithm proposed in `python`'s `scipy.optimize` module. Non-differentiability of the ℓ_1-norm in the elastic-net penalization is dealt with the standard trick of reformulating the problem, using the fact that $|a| = a_+ + a_-$ for $a \in \mathbb{R}$.

2SGD This is Algorithm 1 without variance reduction, namely basic doubly stochastic gradient descent.

2SAG A doubly stochastic version of Stochastic Averaged Gradient from Schmidt et al. (2013) for which every randomly picked gradient ∇f_i is estimated via MCMC.

2SVRG Doubly Stochastic Proximal Gradient descent described in Algorithm 1 above.

SVRG-MB Mini-Batch Prox-SVRG described in Nitanda (2014), which can be seen as Algorithm 3 (see below) with $K_S = 0$. This is a *simply* stochastic algorithm, since there is no MCMC approximation of the gradients ∇f_i. The question of mini-batch sizing is critical and is adressed in Appendix E. We used $n_{\text{mb}} = 0.1n$ or $n_{\text{mb}} = 0.01n$ in our experiments.

The "simply stochastic" counterparts SGD, SAG and SVRG are way slower than the corresponding doubly stochatic versions, since they rely on many computations of stochastic gradients ∇f_i, which are numerically costly, as explained above. Therefore, these algorithms aren't included in our experiments. The same settings are used throughout all experiments, some of them being tuned by hand: steps size for 2SGD, 2SVRG and 2SAG are taken as $\gamma_t = \alpha_0/\sqrt{1 + t/t_0}$ where α_0 and t_0 depend on the dataset, the phase length m for 2SVRG is equal to the number n of failures of each datasets as

suggested in Konečný and Richtárik (2013). Even if the theory allows us to take a constant step-size γ for 2SVRG, the choice of a decreasing one helps curbing the variance introduced by Monte Carlo approximation, and gives good performance. All doubly stochastic algorithms use the same MCMC approximation code, given by Algorithm 2, in the particular case of the Cox partial likelihood.

The number of iterations N_k in Algorithm 2 at each phase k exponentially increases as k increases from 1 to K, the total number of phases. We therefore *quickly* reach the phase k for which $N_k > n$. This problem is solved by considering the following hybrid algorithm.

Hybrid SVRG algorithm Since N_k exponentially increases, the 2SVRG's complexity is higher than SVRG's original complexity. However, the algorithm 2SVRG is very efficient during the first phases: we introduce an hybrid solver that begins with 2SVRG and switches to SVRG with mini-batchs (denoted SVRG-MB). Mini-batching simply consists in replacing single stochastic gradients ∇f_i by an average over a subset $\mathcal{B}$ of size n_{mb} uniformly selected at random. This is useful in our case, since we can use a computational trick (recurrence formula) to compute mini-batched gradients, see Appendix B. In our experiments, we used $n_{\text{mb}} = 0.1n$ or $n_{\text{mb}} = 0.01n$, a constant step-size γ designed for each dataset, and switched from 2SVRG to SVRG-MB after $K_S = 5$ phases. We set $N_k = n^{k/(K_S+2)}$ so that N_k never exceeds n.

Algorithm 3 Hybrid SVRG (HSVRG)

Require: Number of phases before switching $K_S \geq 1$, total number of phases $K \geq K_S$, phase-length $m \geq 1$, step-size $\gamma > 0$, MCMC number of iterations per phase $(N_k)_{k=1}^K$, starting point $\theta^0 \in \mathbb{R}^d$

Initialize: $\tilde{\theta} \leftarrow \theta^0$ and compute $\nabla f_i(\tilde{\theta})$ for $i = 1, \ldots, n$
Apply: Algorithm 1 with K_S phases, phase-length m, step-size γ, MCMC number of iterations per phase $(N_k)_{k=1}^{K_S}$, starting point $\theta^0 \in \mathbb{R}^d$. Outputs $\tilde{\theta}^{K_S}$
for $k = K_S + 1$ **to** K **do**
  for $t = 0$ **to** $m_{\text{mb}} - 1 = \lfloor (m-1)/n_{\text{mb}} \rfloor$ **do**
    Pick a set of random indices $\mathcal{B} \sim (\mathcal{U}[n])^{n_{\text{mb}}}$
    $d^t = \nabla f_{\mathcal{B}}(\theta^t) - \nabla f_{\mathcal{B}}(\tilde{\theta}) + \frac{1}{n}\sum_{j=1}^n \nabla f_j(\tilde{\theta})$
    $\omega^{t+1} \leftarrow \theta^t - \gamma d^t$
    $\theta^{t+1} \leftarrow \text{prox}_{\gamma h}(\omega^{t+1})$
  end for
  Update $\tilde{\theta} \leftarrow \frac{1}{m_{\text{mb}}}\sum_{t=1}^{m_{\text{mb}}} \theta^t$, $\theta^0 \leftarrow \tilde{\theta}$, $\tilde{\theta}^k \leftarrow \tilde{\theta}$
end for
Return: $\tilde{\theta}^K$

Datasets We compare algorithms on the following datasets. The first three are standard benchmarks in survival analysis, the fourth one is a large simulated dataset where the number of observations n exceeds the number of features d. This differs from supervised gene expression data: such a large-scale setting happens for longitudinal clinical trials, medical adverse event monitoring and business data minings tasks.

- NKI70 contains survival data for 144 breast cancer patients, 5 clinical covariates and the expressions from 70 gene signatures, see Van De Vijver et al. (2002).
- Luminal contains survival data for 277 patients with breast cancer who received the adjuvant tamoxifen, with 44,928 expressions measurements, see Loi et al. (2007).
- Lymphoma contains 7399 gene expressions data for 240 lymphoma patients. The data was originally published in Alizadeh et al. (2000).
- We generated a Gaussian features matrix X with $n = 10,000$ observations and $d = 500$ predictors, with a Toeplitz covariance and correlation equal to 0.5. The failure times follow a Weibull distribution. See Appendix D for details on simulation in this model.

We compare in Figures 2 and 3 all algorithms for ridge penalization, namely $\alpha = 0$ and $\lambda = 1/\sqrt{n}$. Experiences with other values of α and λ are given in Appendix (including the Lasso penalization for instance).

Conclusions. The experiments first show that the solvers 2SAG and 2SVRG give better results than 2SGD. Moreover, 2SAG is slightly better than 2SVRG: the descent direction being an average, it averages the previously computed MCMC and then curbs the variance that MCMC induces. However, they less good perform after some phases. Hybrid SVRG enjoys the best of both worlds - it takes both advantage of the 2SVRG's aggressive start and the good asymptotic behavior of SVRG-MB - and gives better performance than the other solvers in most cases.

6 Conclusion

We have proposed a *doubly* stochastic gradient algorithm to extend SGD-like algorithms beyond empirical risk minimization setting. The algorithm we proposed is the result of two different ideas: sampling from uniform distribution to avoid computing a large sum, and sampling using MCMC methods to avoid computing a more complicated expectation. We have also provided theoretical guarantees of convergence for both the convex and the strongly-convex setting.

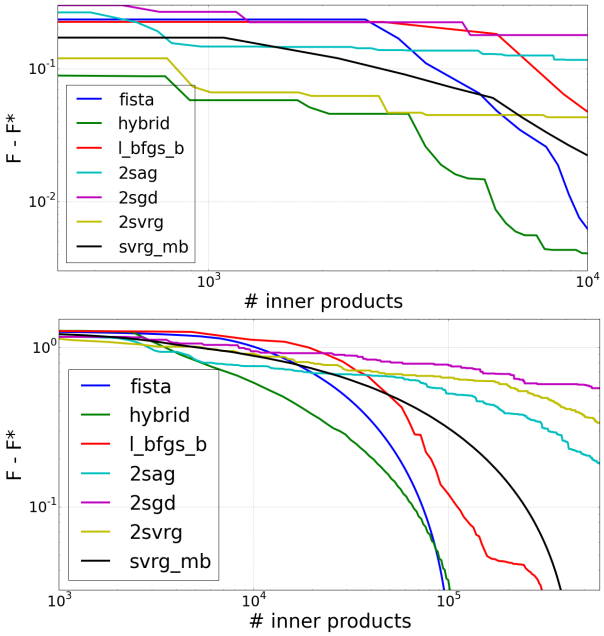

Figure 2: Distance to optimum of all algorithms on NKI70 (top) and Lymphoma (bottom) with ridge penalization ($\alpha = 0$ and $\lambda = 1/\sqrt{n}$)

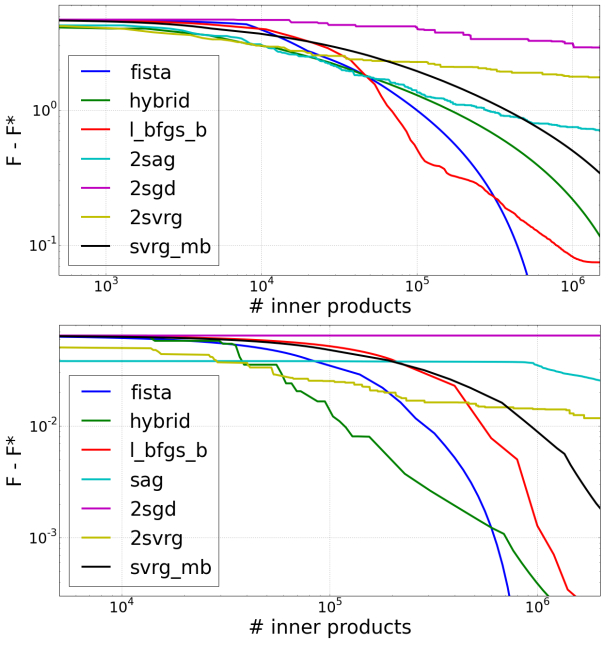

Figure 3: Distance to optimum of all algorithms on Luminal (top) and on the simulated dataset (bottom) with ridge penalization ($\alpha = 0$ and $\lambda = 1/\sqrt{n}$)

This *doubly* stochastic gradient algorithm is very efficient during the early phases. The hybrid version of our algorithm, at the crossing of *simply* and *doubly* stochastic gradient algorithms, significantly outperforms state-of-the-art methods.

As future works, we intend to pursue in two different directions. First, our algorithm can be extended to Conditional Random Fields (CRFs) since each subfunction's gradient takes the form

$$\nabla f_i(\theta) = \nabla(-\log(p(y_i|x_i,\theta)) = \mathbb{E}^{H_i(\theta)\sim\pi^i_\theta}[H_i(\theta)],$$

for a certain random variable $H_i(\theta)$ (see Page 2 in Schmidt et al. (2015)). Plus, even if Algorithm 2 gives enough theoretical guarantees to prove the theorems, we believe that a better design of the MCMC algorithm can improve the practical results: by running many MCMC in parallel and averaging them, and/or by using adaptative MCMC to choose a better proposal distribution.

References

A. A. Alizadeh, M. B. Eisen, R. E. Davis, C. Ma, I. S. Lossos, A. Rosenwald, J. C. Boldrick, H. Sabet, T. Tran, X. Yu, et al. Distinct types of diffuse large b-cell lymphoma identified by gene expression profiling. *Nature*, 403(6769):503–511, 2000.

Y. F. Atchade, G. Fort, and E. Moulines. On stochastic proximal gradient algorithms. *ArXiv e-prints*, February 2014.

F. Bach and E. Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate o (1/n). In *Advances in Neural Information Processing Systems*, pages 773–781, 2013.

A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2(1):183–202, March 2009.

S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. ISBN 0521833787.

D. R. Cox. Regression models and life tables (with discussion). *Journal of the Royal Statistical Society, Series B*, 74:187–220, 1972.

B. Dai, B. Xie, N. He, Y. Liang, A. Raj, M. F. Balcan, and L. Song. Scalable kernel methods via doubly stochastic gradients. *CoRR*, abs/1407.5599, 2014.

A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. *CoRR*, abs/1407.0202, 2014.

J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.*, 12:2121–2159, July 2011.

L. Einav and J. Levin. Economics in the age of big data. *Science*, 346(6210):1243089, 2014.

G. Fort and E. Moulines. Convergence of the monte carlo expectation maximization for curved exponential families. *Ann. Statist.*, 31(4):1220–1259, 08 2003.

J. Goeman. L1 penalized estimation in the ox proportional hazards model. *Biometrical Journal*, 52(1):70–84, 2010.

C. Hu, W. Pan, and J. T. Kwok. Accelerated gradient methods for stochastic optimization and online learning. *Advances in Neural Information Processing Systems*, 2009.

R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems 26*, pages 315–323, 2013.

A. Juditsky and A. S. Nemirovski. First Order Methods for Nonsmooth Convex Large-Scale Optimization, II: Utilizing Problem's Structure. In *Optimization for Machine Learning*, pages 29–63. MIT Press, August 2010.

J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *Ann. Math. Statist.*, 23(3):462–466, 09 1952.

J. Konečný and P. Richtárik. Semi-Stochastic Gradient Descent Methods. *ArXiv e-prints*, December 2013.

J. D. Lafferty, A. McCallum, and F. C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 282–289, 2001. ISBN 1-55860-778-1.

G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming Series A*, 2010.

D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *Math. Program.*, 45(3):503–528, December 1989.

S. Loi, B. Haibe-Kains, C. Desmedt, F. Lallemand, A. M. Tutt, C. Gillet, P. Ellis, A. Harris, J. Bergh, J. A. Foekens, et al. Definition of clinically distinct molecular subtypes in estrogen receptor–positive breast carcinomas through genomic grade. *Journal of clinical oncology*, 25(10):1239–1246, 2007.

S. Mittal, D. Madigan, J. Cheng, and R. S. Burd. Large-scale parametric survival analysis. *Statistics in medicine*, 32(23):3955–3971, 10 2013.

T. B. Murdoch and A. S. Detsky. The inevitable application of big data to health care. *Jama*, 309(13):1351–1352, 2013.

A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4): 1574–1609, 2009.

M. Neumaier. How big data is transforming the competitive landscape of business. `http://actuaries.asn.au/Library/Events/FSF/2014/Neumaier4a.pdf`, 2014. Accessed: 2010-09-30.

A. Nitanda. Stochastic proximal gradient descent with acceleration techniques. *Advances in Neural Information Processing Systems*, pages 1574–1582, 2014.

M. Y. Park and T. Hastie. L1-regularization path algorithm for generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):659–677, 2007.

H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 09 1951.

C. P. Robert and G. Casella. *Monte Carlo statistical methods*, volume 319. Springer New York, 2004.

M. Schmidt, N. Le Roux, and F. Bach. Minimizing Finite Sums with the Stochastic Average Gradient. *ArXiv e-prints*, September 2013.

M. Schmidt, R. Babanezhad, M. Osama Ahmed, A. Defazio, A. Clifton, and A. Sarkar. Non-Uniform Stochastic Average Gradient Method for Training Conditional Random Fields. *ArXiv e-prints*, April 2015.

S. Shalev-Shwartz and T. Zhang. Proximal Stochastic Dual Coordinate Ascent. *ArXiv e-prints*, November 2012.

N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for cox's proportional hazards model via coordinate descent. *Journal of Statistical Software*, 39(5):1–13, 3 2011.

I. Sohn, J. Kim, S.-H. Jung, and C. Park. Gradient lasso for cox proportional hazards model. *Bioinformatics*, 25(14):1775–1781, 2009.

R. Tibshirani. The lasso method for variable selection in the cox model. *Statistics in Medicine*, 16:385–395, 1997.

R. Tibshirani and T. Hastie. *Generalized Additive Models*, chapter 8. Chapman and Hall, London, 1990.

P. Toulis and E. M. Airoldi. Implicit stochastic gradient descent for principled estimation with large datasets. *ArXiv e-prints*, August 2014.

M. J. Van De Vijver, Y. D. He, L. J. van't Veer, H. Dai, A. A. M. Hart, D. W. Voskuil, G. J. Schreiber, J. L. Peterse, C. Roberts, M. J. Marton, et al. A gene-expression signature as a predictor of survival in breast

cancer. *New England Journal of Medicine*, 347(25): 1999–2009, 2002.

V. N. Vapnik. An overview of statistical learning theory. *Neural Networks, IEEE Transactions on*, 10(5):988–999, 1999.

D. M. Witten and R. Tibshirani. Survival analysis with high-dimensional covariates. *Statistical methods in medical research*, 2009.

L. Xiao and T. Zhang. A Proximal Stochastic Gradient Method with Progressive Variance Reduction. *ArXiv e-prints*, March 2014.

X. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. In *In Advances in Neural Information Processing Systems 23*, 2009.

Y. Yang and H. Zhou. A cocktail algorithm for solving the elastic net penalized cox's regression in high dimensions. *Statistics and Its Interface*, 6:167–173, 2013.

T. Zhang and F. Oles. The value of unlabeled data for classification problems. In *Proceedings of the Seventeenth International Conference on Machine Learning,(Langley, P., ed.)*, pages 1191–1198. Citeseer, 2000.

A Proofs

A.1 Proof of Proposition 4.3

We first prove Proposition 4.3 that ensures that Algorithm 2 provides the bounds of Assumption 2.

Proof. Since there exists $M > 0$ such that the proposal Q and the stationary distribution π satisfy $\pi(x) \leq MQ(x)$, for all x in the support of π, the Theorem 7.8 in Robert and Casella (2004) states that the Algorithm 2 produces a geometrically ergodic Markov kernel P with ergodicity constants uniformly controlled:

$$\|P^k(x,\cdot) - \pi\|_{TV} \leq 2\left(1 - \frac{1}{M}\right)^k, \tag{8}$$

where P^k is the kernel of the k^{th} iteration of the algorithm and $\|\cdot\|_{TV}$ is the total variation norm. Since $\widehat{\nabla} f_{i_t}(\theta^{t-1})$ is computed as the mean of the iterates of the Markov chain, a simple computation enables us to bound the bias of the error and Proposition 12 from Fort and Moulines (2003) gives the upper bound for the expected squared error:

$$\|\mathbb{E}_t \eta^t\| \leq \frac{C_1}{N_k} \text{ and } \mathbb{E}_t\|\eta^t\|^2 \leq \frac{C_2}{N_k} \tag{9}$$

where C_1 and C_2 are some finite constants, and N_k the number of iterations of the Markov chain. It can be shown that $C_1 = 2M$ and that C_2 is related to a constant from the Rosenthal's inequality. ■

A.2 Preliminaries to the proofs of Theorems 4.1 and 4.2

In what follows, the key lemmas for the proofs of Theorems 4.1 and 4.2 are stated and proved when not directly borrowed from previous articles.

Lemma A.1. *For $\Delta^t := \widehat{\nabla} f_{i_t}(\theta^{t-1}) - \nabla f_{i_t}(\tilde{\theta}) + \nabla f(\tilde{\theta}) - \nabla f(\theta^{t-1})$, we have:*

$$\begin{aligned}\mathbb{E}_t\|\Delta^t\|^2 \leq 8L[&F(\theta^{t-1}) - F(\theta^*) \\ &+ F(\tilde{\theta}) - F(\theta^*)] + 3\mathbb{E}_t\|\eta^t\|^2.\end{aligned}$$

The proof of Lemma A.1 uses Lemma 1 in Xiao and Zhang (2014).

Lemma A.2. *Johnson and Zhang (2013); Xiao and Zhang (2014) Consider F satisfying Assumption 1. Then,*

$$\frac{1}{n}\sum_{i=1}^n \|\nabla f_i(\theta) - \nabla f_i(\theta^*)\|^2 \leq 2L[F(\theta) - F(\theta^*)]$$

Proof of Lemma A.1. For the sake of simplicity, we now denote $d_i^t = \nabla f_i(\theta^{t-1}) - \nabla f_i(\tilde{\theta})$ and $d^t = \nabla f(\theta^{t-1}) - \nabla f(\tilde{\theta})$, so that one gets $\Delta^t = d_{i_t}^t - d^t + \eta^t$. Then, using the expectation introduced in Section 4, we repeatedly use the identity $\mathbb{E}_t\|\xi\|^2 = \mathbb{E}_t\|\xi - \mathbb{E}_t\xi\|^2 + \|\mathbb{E}_t\xi\|^2$. First with $\xi = \Delta^t$ (since $\mathbb{E}_t d_{i_t}^t = d^t$, one gets $\mathbb{E}_t\xi = \mathbb{E}_t\eta^t$) :

$$\mathbb{E}_t\|\Delta^t\|^2 = \mathbb{E}_t\|d_{i_t}^t + \eta^t - \left(d^t + \mathbb{E}_t\eta^t\right)\|^2 + \|\mathbb{E}_t\eta^t\|^2$$

then, successively with $\xi = d_{i_t}^t + \eta^t$, $\xi = d^t + \eta^t$ and finally $\xi = \eta^t$:

$$\begin{aligned}\mathbb{E}_t\|\Delta^t\|^2 &= \mathbb{E}_t\|d_{i_t}^t + \eta^t\|^2 + \|\mathbb{E}_t\eta^t\|^2 - \|d^t + \mathbb{E}_t\eta^t\|^2 \\ &= \mathbb{E}_t\|d_{i_t}^t + \eta^t\|^2 + \|\mathbb{E}_t\eta^t\|^2 \\ &\quad - \left(\mathbb{E}_t\|d^t + \eta^t\|^2 - \mathbb{E}_t\|\eta^t - \mathbb{E}_t\eta^t\|^2\right) \\ &= \mathbb{E}_t\|d_{i_t}^t + \eta^t\|^2 + \mathbb{E}_t\|\eta^t\|^2 - \mathbb{E}_t\|d^t + \eta^t\|^2.\end{aligned}$$

Now we remark that $\mathbb{E}_t\|d^t + \eta^t\|^2 \geq 0$, and the identity $\|a+b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ gives the majoration

$$\mathbb{E}_t\|\Delta^t\|^2 \leq 2\mathbb{E}_t\|d_{i_t}^t\|^2 + 3\mathbb{E}_t\|\eta^t\|^2.$$

Now rewriting $d_{i_t}^t = \nabla f_{i_t}(\theta^{t-1}) - \nabla f_{i_t}(\theta^*) + \nabla f_{i_t}(\theta^*) - \nabla f_{i_t}(\tilde{\theta})$, the same identity leads to

$$\begin{aligned}\mathbb{E}_t\|\Delta^t\|^2 \leq\ &4\mathbb{E}_t\|\nabla f_{i_t}(\theta^{t-1}) - \nabla f_{i_t}(\theta^*)\|^2 \\ &+ 4\mathbb{E}_t\|\nabla f_{i_t}(\tilde{\theta}) - \nabla f_{i_t}(\theta^*)\|^2 + 3\mathbb{E}_t\|\eta^t\|^2.\end{aligned}$$

The desired result follows applying twice Lemma A.2. ■

When F is μ-strongly convex, the next Lemma (Lemma 3 in Xiao and Zhang (2014)) provides a key lower bound.

Lemma A.3. *Xiao and Zhang (2014) Consider $F = f + h$ satifying Assumption 1, where f is L_f-smooth, $L_f > 0$, f is μ_f-strongly convex, $\mu_f \geq 0$, h is μ_h-strongly convex, $\mu_h \geq 0$. For any $x, v \in \mathbb{R}^d$, we define $x^+ = \mathrm{prox}_{\gamma h}(x - \gamma v)$, $g = \frac{1}{\gamma}(x - x^+)$, where $\gamma \in (0, \frac{1}{L_f}]$. Then, for any $y \in \mathbb{R}^d$:*

$$\begin{aligned}F(y) \geq F(x^+) + g^\top(y - x) + \frac{\gamma}{2}\|g\|^2 + \frac{\mu_f}{2}\|y - x\|^2 \\ + \frac{\mu_h}{2}\|y - x^+\|^2 + (v - \nabla f(x))^\top(x^+ - y).\end{aligned} \tag{10}$$

Remark 3. Note that in Lemma A.3, one can freely choose μ_f and μ_h (in particular one can take $\mu_f = 0$ or $\mu_h = 0$), as long as $\mu_f + \mu_h = \mu$.

The following Lemma comes from Atchade et al. (2014) (Lemma 14):

Lemma A.4. *Atchade et al. (2014) Consider $F = f + h$ satifying Assumption 1, where f is L_f-smooth, and $T_\gamma : x \mapsto \mathrm{prox}_{\gamma h}[x - \gamma\nabla f(x)]$ with $\gamma \in (0, 2/L_f]$. Let $x, y \in \mathbb{R}^d$, we have:*

$$\|T_\gamma(x) - T_\gamma(y)\| \leq \|x - y\|$$

A.3 Proof of Theorem 4.1

Proof. The proof begins with the study of the distance $\|\theta^t - \theta^*\|^2$ *between* the phases $k-1$ and k. To ease the reading, when staying between these two phases, we write $\tilde{\theta}$ instead of $\tilde{\theta}^{k-1}$. Introducing $g^t = \frac{1}{\gamma}(\theta^{t-1} - \theta^t)$, we may write:

$$\begin{aligned}\|\theta^t - \theta^*\|^2 &= \|\theta^{t-1} - \gamma g^t - \theta^*\|^2 \\ &= \|\theta^{t-1} - \theta^*\|^2 - 2\gamma(g^t)^\top(\theta^{t-1} - \theta^*) \\ &\quad + \gamma^2\|g^t\|^2.\end{aligned}$$

To upper bound the term $-2\gamma(g^t)^\top(\theta^{t-1} - \theta^*) + \gamma^2\|g^t\|^2$, we apply the Lemma A.3 with $x = \theta^{t-1}$, $x^+ = \theta^t$ and $y = \theta^*$. With again $\Delta^t = \widehat{\nabla} f_{i_t}(\theta^{t-1}) - \nabla f_{i_t}(\tilde{\theta}) + \nabla f(\tilde{\theta}) - \nabla f(\theta^{t-1})$, we obtain

$$\begin{aligned}&-(g^t)^\top(\theta^{t-1} - \theta^*) + \frac{\gamma}{2}\|g^t\|^2 \\ &\quad\leq F(\theta^*) - F(\theta^t) - \frac{\mu_f}{2}\|\theta^{t-1} - \theta^*\|^2 \\ &\quad - \frac{\mu_h}{2}\|x^t - \theta^*\|^2 - (\Delta^t)^\top(\theta^t - \theta^*),\end{aligned}$$

and

$$\begin{aligned}\|\theta^t - \theta^*\|^2 \leq \|\theta^{t-1} - \theta^*\|^2 + 2\gamma[F(\theta^*) - F(\theta^t)] \\ - 2\gamma(\Delta^t)^\top(\theta^t - \theta^*). \qquad (11)\end{aligned}$$

We now concentrate on the quantity $-2\gamma(\Delta^t)^\top(\theta^t - \theta^*)$. Introducing $\nu^t = \text{prox}_{\gamma h}[\theta^{t-t1} - \gamma\nabla f(\theta^{t-1})] \in \mathcal{F}_{t-1}$ i.e. the vector obtained from θ^{t-1} after an exact proximal gradient descent step, we get

$$\begin{aligned}&-2\gamma(\Delta^t)^\top(\theta^t - \theta^*) \\ &\quad= -2\gamma(\Delta^t)^\top(\theta^t - \nu^t) - 2\gamma(\Delta^t)^\top(\nu^t - \theta^*) \\ &\quad\leq 2\gamma\|\Delta^t\| \cdot \|\theta^t - \nu^t\| - 2\gamma(\Delta^t)^\top(\nu^t - \theta^*)\end{aligned}$$

where the inequality follows from the Cauchy-Schwartz inequality. Now the non-expansiveness property of proximal operators $\|\text{prox}_{\gamma h}(x) - \text{prox}_{\gamma h}(y)\| \leq \|x - y\|$ leads to

$$\begin{aligned}&-2\gamma(\Delta^t)^\top(\theta^t - \theta^*) \\ &\quad\leq 2\gamma\|\Delta^t\| \cdot \|\{\theta^{t-1} - \gamma(\Delta^t + \nabla f(\theta^{t-1}))\} \\ &\quad - \{\theta^{t-1} - \gamma\nabla f(\theta^{t-1})\}\| - 2\gamma(\Delta^t)^\top(\nu^t - \theta^*) \\ &\quad\leq 2\gamma^2\|\Delta^t\|^2 - 2\gamma(\Delta^t)^\top(\nu^t - \theta^*).\end{aligned}$$

Reminding that $\nu^t \in \mathcal{F}_{t-1}$, we derive:

$$\begin{aligned}&-2\gamma\mathbb{E}_t(\Delta^t)^\top(\theta^t - \theta^*) \\ &\quad\leq 2\gamma^2\mathbb{E}_t\|\Delta^t\|^2 - 2\gamma(\mathbb{E}_t\Delta^t)^\top(\nu^t - \theta^*) \\ &\quad\leq 2\gamma^2\mathbb{E}_t\|\Delta^t\|^2 + 2\gamma\|\mathbb{E}_t\Delta^t\| \cdot \|\nu^t - \theta^*\|,\end{aligned}$$

the last inequality comes from the Cauchy-Schwartz inequality. Since θ^* is the minimum of $F = f + h$, it satisfies $\theta^* = \text{prox}_{\gamma h}[\theta^* - \gamma\nabla f(\theta^*)]$. Thus, the Lemma A.4 and the Assumption 1 on the sequence (θ^t) give us $\|\nu^t - \theta^*\| \leq \|\theta^{t-1} - \theta^*\| \leq B$. We also remark that $\mathbb{E}_t\Delta^t = \mathbb{E}_t\eta^t$. For all t between phases $k-1$ and k, we finally apply Lemma A.1 to obtain:

$$\begin{aligned}&-2\gamma\mathbb{E}_t(\Delta^t)^\top(\theta^t - \theta^*) \\ &\quad\leq 16\gamma^2 L[F(\theta^{t-1}) - F(\theta^*) + F(\tilde{\theta}) - F(\theta^*)] \\ &\quad + 6\gamma^2\mathbb{E}_t\|\eta^t\|^2 + 2\gamma B\|\mathbb{E}_t\eta^t\|. \qquad (12)\end{aligned}$$

Taking the expectation $\mathbb{E}_t$ on inequation (11) and combining with previous inequality leads to

$$\begin{aligned}\mathbb{E}_t\|\theta^t - \theta^*\|^2 &\leq \|\theta^{t-1} - \theta^*\|^2 + 2\gamma[F(\theta^*) - F(\theta^t)] \\ &\quad + 16\gamma^2 L[F(\theta^{t-1}) - F(\theta^*) + F(\tilde{\theta}) - F(\theta^*)] \\ &\quad + 6\gamma^2\mathbb{E}_t\|\eta^t\|^2 + 2\gamma B\|\mathbb{E}_t\eta^t\|.\end{aligned}$$

With the notation of Algorithm 1, $\tilde{\theta} = \tilde{\theta}^{k-1} = \theta^0$. Now, applying iteratively the previous inequality over $t = 1, 2, \ldots, m$ and taking the expectation $\mathbb{E}$ over $i_1, \theta^1, i_2, \theta^2, \ldots, i_m, \theta^m$, we obtain:

$$\begin{aligned}&\mathbb{E}\|\theta^m - \theta^*\|^2 + 2\gamma[\mathbb{E}F(\theta^m) - F(\theta^*)] \\ &+2\gamma(1 - 8L\gamma)\sum_{t=1}^{m-1}[\mathbb{E}F(\theta^t) - F(\theta^*)] \\ &\quad\leq \|\theta^0 - \theta^*\|^2 + 16L\gamma^2[F(\theta^0) - F(\theta^*) + m(F(\tilde{\theta}) \\ &\quad - F(\theta^*))] + 6\gamma^2\sum_{t=1}^{m}\mathbb{E}\|\eta^t\|^2 + 2\gamma B\sum_{t=1}^{m}\mathbb{E}\|\mathbb{E}_t\eta^t\|.\end{aligned}$$

Now, by convexity of F and the definition $\tilde{\theta}^k = \frac{1}{m}\sum_{t=1}^{m}\theta^t$, we may write $F(\tilde{\theta}^k) \leq \frac{1}{m}\sum_{t=1}^{m}F(\theta^t)$. Noticing that $2\gamma(1 - 8L\gamma) < 2\gamma$ leads to

$$\begin{aligned}&2\gamma(1 - 8L\gamma)m[\mathbb{E}F(\tilde{\theta}^k) - F(\theta^*)] \\ &\quad\leq \|\tilde{\theta} - \theta^*\|^2 + 16L\gamma^2(m+1)[F(\tilde{\theta}) - F(\theta^*)] \\ &\quad + 6\gamma^2\sum_{t=1}^{m}\mathbb{E}\|\eta^t\|^2 + 2\gamma B\sum_{t=1}^{m}\|\mathbb{E}\eta^t\|.\end{aligned}$$

Under the Assumption 2, we have

$$\begin{aligned}&6\gamma^2\sum_{t=1}^{m}\mathbb{E}\|\eta^t\|^2 + 2\gamma B\sum_{t=1}^{m}\|\mathbb{E}\eta^t\| \\ &\qquad\leq (6\gamma^2 C_2 + 2\gamma B C_1)\frac{m}{N_k}\end{aligned}$$

whereas the μ-strong convexity of F implies $\|\tilde{\theta}^{k-1} - \theta^*\|^2 \leq \frac{2}{\mu}[F(\tilde{\theta}^{k-1}) - F(\theta^*)]$. This leads to

$$\mathbb{E}F(\tilde{\theta}^k) - F(\theta^*) \leq \rho\left(\mathbb{E}F(\tilde{\theta}^{k-1}) - F(\theta^*)\right) + \frac{D}{N_k}$$

for D and ρ as defined in the theorem. Applying the last inequality recursively leads to the result. ■

A.4 Proof of Theorem 4.2

Proof. As at the begining of the proof of Theorem 4.1, we consider that we stand between phase $k-1$ and phase k of Algorithm 1 and consequently $\theta^0 = \tilde{\theta}^{k-1}$. We use the same arguments until (11), with the difference that, in this non-strongly convex case, we have $\mu_f = \mu_h = 0$. We obtain for all t between phases 1 and m

$$F(\theta^t) - F(\theta^*) \leq \frac{1}{2\gamma}(\|\theta^{t-1} - \theta^*\|^2 - \|\theta^t - \theta^*\|^2) - (\theta^t - \theta^*)^\top \Delta^t.$$

Summing over $t = 1, \dots, \tau$ (for $\tau \leq m$) leads to

$$\sum_{t=1}^{\tau}[F(\theta^t) - F(\theta^*)] \leq \frac{1}{2\gamma}(\sum_{t=0}^{\tau-1}\|\theta^t - \theta^*\|^2 - \sum_{t=1}^{\tau}\|\theta^t - \theta^*\|^2) - \sum_{t=1}^{\tau}(\theta^t - \theta^*)^\top \Delta^t. \tag{13}$$

We now use Equation (13) (with $\tau = m$) and the convexity of $\|\cdot\|^2$ with $\tilde{\theta}^k = \frac{1}{m}\sum_{t=1}^m \theta^t$ to write

$$\sum_{t=1}^{m}[F(\theta^t) - F(\theta^*)] \leq \frac{1}{2\gamma}\left(\sum_{t=0}^{m-1}\|\theta^t - \theta^*\|^2 - m\|\tilde{\theta}^k - \theta^*\|^2\right) - \sum_{t=1}^{m}(\theta^t - \theta^*)^\top \Delta^t. \tag{14}$$

Starting from Equation (13) again but now summing over $l = 1, \dots, t$, we get

$$\frac{1}{2\gamma}(\|\theta^0 - \theta^*\|^2 - \|\theta^t - \theta^*\|^2) - \sum_{l=1}^{t}(\theta^l - \theta^*)^\top \Delta^l \geq \sum_{l=1}^{t}[F(\theta^l) - F(\theta^*)] \geq 0, \tag{15}$$

where the last inequality follows from the definition of θ^*. In (14), we now substitute $\|\theta^t - \theta^*\|^2$ by the upper bound derived from (15) to write (noticing that $\theta^0 = \tilde{\theta}^{k-1}$):

$$\begin{aligned}
&\sum_{t=1}^{m}[F(\theta^t) - F(\theta^*)] \\
&\leq \frac{m}{2\gamma}(\|\tilde{\theta}^{k-1} - \theta^*\|^2 - \|\tilde{\theta}^k - \theta^*\|^2) - \sum_{t=1}^{m-1}\sum_{l=1}^{t}(\theta^l - \theta^*)^\top \Delta^l - \sum_{t=1}^{m}(\theta^t - \theta^*)^\top \Delta^t \\
&\leq \frac{m}{2\gamma}(\|\tilde{\theta}^{k-1} - \theta^*\|^2 - \|\tilde{\theta}^k - \theta^*\|^2) - \sum_{t=1}^{m}(m+1-t)(\theta^t - \theta^*)^\top \Delta^t.
\end{aligned}$$

As in the proof of Theorem 4.1 (see Equation (12)), each term $-\mathbb{E}_t(\theta^t - \theta^*)^\top \Delta^t$ is upper bounded by $8\gamma L[F(\theta^{t-1}) - F(\theta^*) + F(\tilde{\theta}^{k-1}) - F(\theta^*)] + 3\gamma\mathbb{E}_t||\eta^t||^2 + B||\mathbb{E}_t\eta^t||$. Now with $m+1-t \leq m$ and Assumption 2, we obtain:

$$\begin{aligned}
&\frac{1}{m}\sum_{t=1}^{m}\mathbb{E}[F(\theta^t) - F(\theta^*)] \\
&\leq \frac{1}{2\gamma}(\|\tilde{\theta}^{k-1} - \theta^*\|^2 - \mathbb{E}\|\tilde{\theta}^k - \theta^*\|^2) \\
&+ 8L\gamma\Big\{\sum_{t=1}^{m}[\mathbb{E}F(\theta^{t-1}) - F(\theta^*)] + F(\theta^*) - \mathbb{E}[F(\theta^m)] \\
&+ (m+1)[\mathbb{E}[F(\tilde{\theta}^{k-1})] - F(\theta^*)]\Big\} + m\frac{3\gamma C_2 + BC_1}{N_k}.
\end{aligned}$$

By definition of γ, we have $8Lm\gamma < 1$, and we can use the convexity of F to lower bound the left hand side. With the inequality $\mathbb{E}[F(\theta^m)] - F(\theta^*) \geq 0$, one has:

$$\begin{aligned}
&(1 - 8L\gamma m)\left[\mathbb{E}[F(\tilde{\theta}^k)] - F(\theta^*)\right] \\
&\leq \frac{1}{2\gamma}\left(\|\tilde{\theta}^{k-1} - \theta^*\|^2 - \mathbb{E}\|\tilde{\theta}^k - \theta^*\|^2\right) \\
&+ 8L\gamma(m+1)\left[\mathbb{E}[F(\tilde{\theta}^{k-1})] - F(\theta^*)\right] \\
&+ m\frac{3\gamma C_2 + BC_1}{N_k}
\end{aligned}$$

We now take the expectation $\mathbb{E}$ on all iterates of the algorithm i.e. on the iterates $i_1, \theta^1, i_2, \theta^2, \dots, i_m, \theta^m$ from the first phase. Introduce the notations $A^k = \mathbb{E}[F(\tilde{\theta}^k)] - F(\theta^*)$ and $a = (8L\gamma(m+1))/(1 - 8Lm\gamma) < 1$, last inequality leads to:

$$\begin{aligned}
&A^k - aA^{k-1} \\
&\leq \frac{1}{2\gamma(1 - 8Lm\gamma)}\left(\mathbb{E}\|\tilde{\theta}^{k-1} - \theta^*\|^2 - \mathbb{E}\|\tilde{\theta}^k - \theta^*\|^2\right) \\
&+ \frac{D}{N_k},
\end{aligned}$$

where D is defined in the theorem. Summing over the phases $k = 1, 2, \dots, K+1$ and lower bounding A^{K+1} with 0, we obtain:

$$(1-a)\sum_{k=1}^{K} A^k \leq aA^0 + \frac{1}{2\gamma(1-8Lm\gamma)}\|\tilde{\theta}^0 - \theta^*\|^2 + \sum_{k=1}^{K+1}\frac{D}{N_k}$$

The last argument is the use of the convexity of F. Remark the explicit forms of the constants in the theorem:

$$D_1 = \frac{a}{1-a}A^0 + \frac{1}{1-a}\frac{\|\tilde{\theta}^0 - \theta^*\|^2}{2\gamma(1-8Lm\gamma)}$$

and $D_2 = \frac{D}{1-a}$. ■

B SGD-like algorithms don't suit well our setting

Many supervised learning problems can be expressed as the minimization of a finite-sum structured function $F(\theta)$. This minimization problem comes either from maximum likelihood estimation (MLE) or from empirical risk minimization (ERM), see Vapnik (1999). In the classical framework of supervised machine learning, given n i.i.d samples $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$, a kernel function $\Phi : \mathcal{X} \to \mathbb{R}^d$ for the prediction $\theta^\top \Phi(x)$ as a linear function of the features, and a loss function $\ell : \mathcal{Y} \times \mathbb{R} \to \mathbb{R}$, the ERM problem writes

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n}\sum_{i=1}^{n} \ell(y_i, \theta^\top \Phi(x_i)) + h(\theta), \tag{16}$$

where $h(\theta)$ is a regularizer. MLE problems are more general: some can be written as 16, while others are complicated. For the Cox partial likelihood, the minimization problem is

$$\min_{\theta \in \mathbb{R}^d} \frac{1}{n}\sum_{i \in D} \ell_i(\theta^\top X_i) + h(\theta), \tag{17}$$

where $X_i = [x_j : j \in R_i] \in \mathbb{R}^{d \times |R_i|}$ and where $\ell_i(z) = -z_i + \log\left(\sum_{j \in R_i} \exp(z_j)\right)$ for $z \in \mathbb{R}^{|R_i|}$.

SGD-like algorithms suit well Problem (16) since computing time of $(1/n)\sum_{i=1}^{n} \nabla_\theta \ell(y_i, \theta^\top \Phi(x_i))$ is obviously n times the computing time of $\nabla_\theta \ell(y_i, \theta^\top \Phi(x_i))$. In our setting, the computation of the gradient $\nabla_\theta f_i(\theta) = \nabla_\theta \ell_i(\theta^\top X_i)$ involves the computation of $|R_i|$ inner products $\theta^\top x_j$, where $|R_i|$ goes from 1 to n_{pat}. However, a computation trick explained below - based on the fact that $(R_i)_{i \in D}$ is a decreasing sequence - enables the computation of the full gradient $\nabla_\theta f(\theta) = (1/n)\sum_{i \in D} \nabla_\theta \ell_i(\theta^\top X_i)$ using only n_{pat} inner products.

Algorithm 4 Computation of $\nabla_\theta f(\theta)$ with n_{pat} inner products

Require: All inner products $(x_j^\top \theta)_{j=1}^{n_{\text{pat}}}$
$S_1 = \sum_{j \in R_{i_n}} x_j \exp(x_j^\top \theta)$
$S_2 = \sum_{j \in R_{i_n}} \exp(x_j^\top \theta)$
$S = S_1/S_2$
for $i = i_{n-1}$ **to** i_1 **do**
  $S_1 = S_1 + \sum_{j \in R_i, j \notin R_{i+1}} x_j \exp(x_j^\top \theta)$
  $S_2 = S_2 + \sum_{j \in R_i, j \notin R_{i+1}} \exp(x_j^\top \theta)$
  $S = S + S_1/S_2$
end for

C Supplementary experiments

We have tested all algorithms with other settings for the penalization. Namely, we considered:

High lasso. We take $\alpha = 1$ and $\lambda = 1/\sqrt{n}$ and illustrate our results in Figure 5.

Low lasso. We take $\alpha = 1$ and $\lambda = 1/n$ and illustrate our results in Figure 4.

High ridge. We take $\alpha = 0$ and $\lambda = 1/\sqrt{n}$ and illustrate our results in Figures 2 and Figures 3.

Low ridge. We take $\alpha = 1$ and $\lambda = 1/n$ and illustrate our results in Figure 6.

D Simulation of data

With Cox model, the hazard ratio for the failure time T_i of the i^{th} patient takes the form:

$$\lambda_i(t) = \lambda_0(t)\exp(x_i^\top \theta),$$

where $\lambda_0(t)$ is a baseline hazard ratio, and $x_i \in \mathbb{R}^d$ the covariates of the i^{th} patient.
We first simulate the feature matrix $X \in \mathbb{R}^{n \times d}$ as a Gaussian vector with a Toepliz covariance, where the correlation between features j and j' is equal to $\rho^{|j-j'|}$, for some $\rho \in (0, 1)$.

We want now to simulate the observed time y_i that corresponds to x_i. We denote the cumulative hazard function $\Lambda(t) = \int_0^t \lambda(s)ds$. Using the definition $\lambda(t) = \frac{f(t)}{1-F(t)}$, we know that $\Lambda(t) = -\log(1 - F(t))$, where f is the p.d.f. and F is the c.d.f. of T.
It is easily seen that $\Lambda(T)$ has distribution $\text{Exp}(1)$ (Exponential with intensity equal to 1): since Λ is an increasing

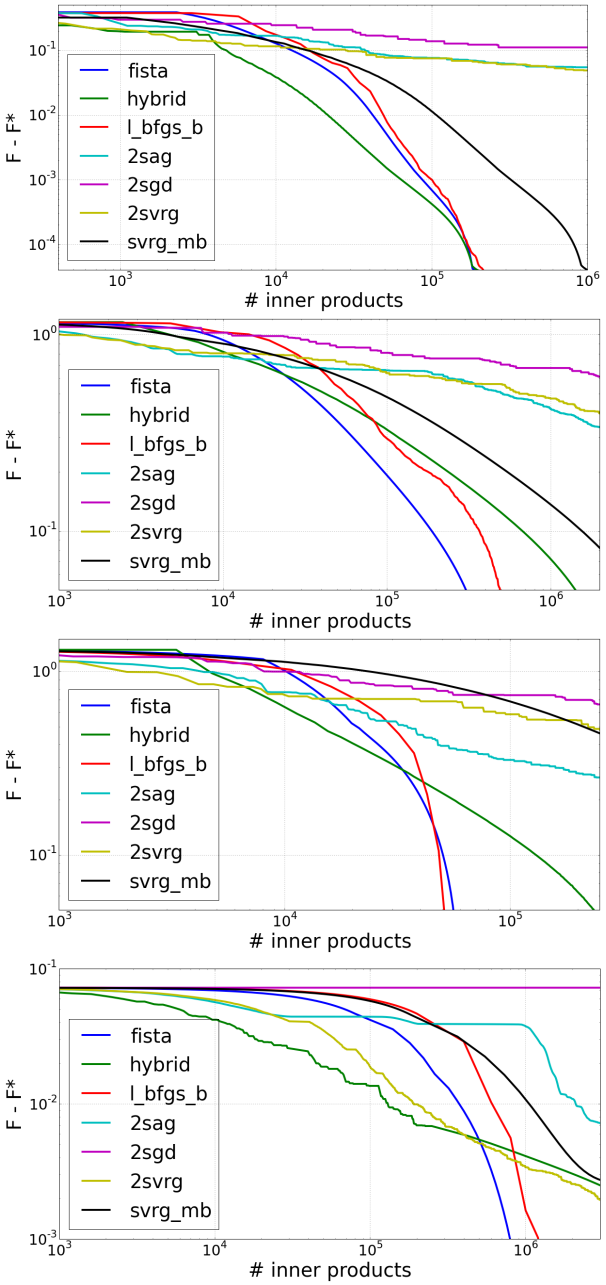

Figure 4: Distance to optimum of all algorithms on NKI70, Lymphoma, Luminal and on the simulated dataset (respectively from top to bottom) for **Low-ridge** penalization

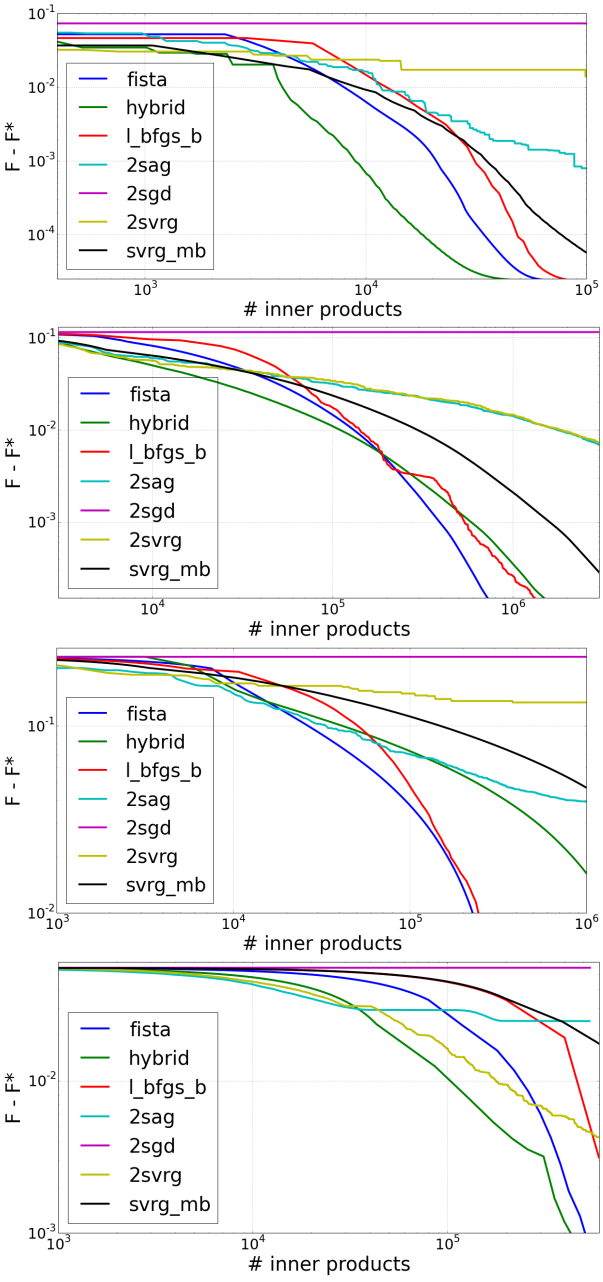

Figure 5: Distance to optimum of all algorithms on NKI70, Lymphoma, Luminal and on the simulated dataset (respectively from top to bottom) for **High-lasso** penalization

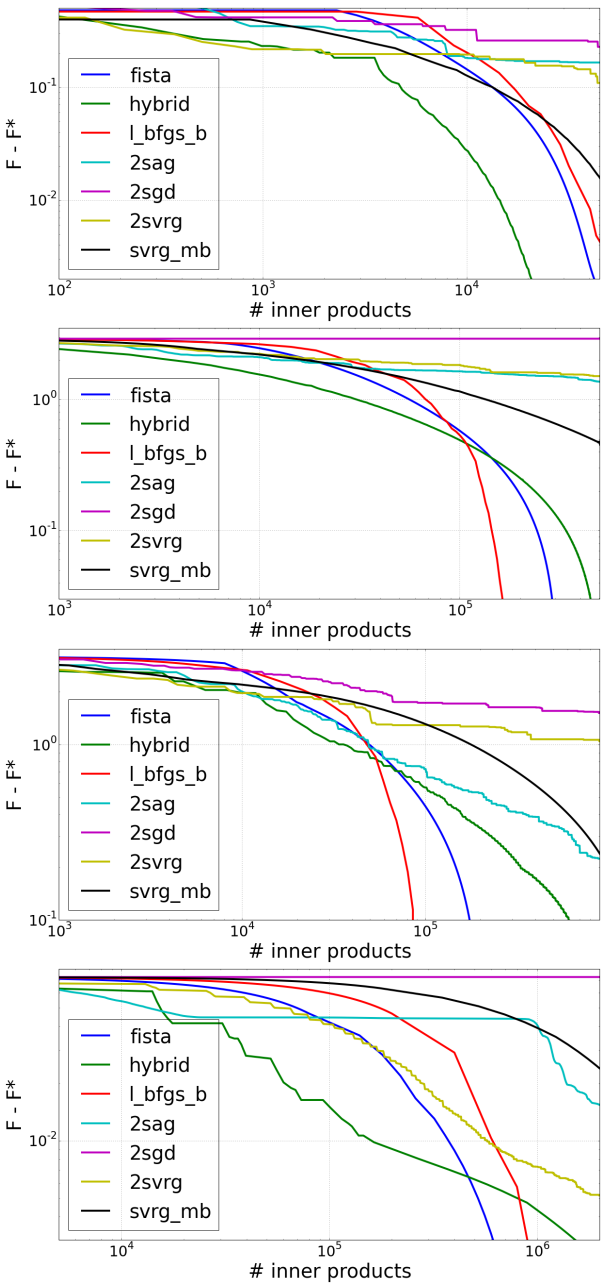

Figure 6: Distance to optimum of all algorithms on NKI70, Lymphoma, Luminal and on the simulated dataset (respectively from top to bottom) for **Low-lasso** penalization

function, we have

$$\begin{aligned}\mathbb{P}(\Lambda(T) \geq t) &= \mathbb{P}(T \geq \Lambda^{-1}(t)) = \int_{\Lambda^{-1}(t)}^{\infty} f(s)ds \\ &= 1 - F(\Lambda^{-1}(t)) \\ &= \exp(-\Lambda(\Lambda^{-1}(t))) \\ &= \exp(-t),\end{aligned}$$

so that simulating failure times is simply achieved by using $T_i = \Lambda^{-1}(E_i)$ where $E_i \sim \text{Exp}(1)$. To compute Λ, we should have a parametric form for λ_0. We assume that T follows the Weibull distribution $\mathcal{W}(1, \nu)$ (when $x_i = 0$). This choice is motivated by the following facts:

- Its cumulative hazard function is easy to invert. Indeed the hazard ratio is given by $\lambda_0(t) = \frac{\nu t^{\nu-1} e^{-t^\nu}}{1-(1-e^{-t^\nu})} = \nu t^{\nu-1}$, so that $\Lambda^{-1}(y) = \left(\frac{y}{\exp(x_i^\top \theta)}\right)^{1/\nu}$.
- It enables two different trends - increasing or decreasing – for the baseline hazard ratio that correspond to two typical behaviours in the medical field.
  - decreasing: after taking a treatment, time before a side-effect's appearence
  - increasing: no memory process and patient's health is worsening

This method enables us to simulate n failures times $T_1, T_2, \ldots, T_n$.
Then, we simulate $C_1, C_2, \ldots, C_n$ with exponential distribution. This finally gives us a set of observed times $(y_i)_{i=1}^n = (T_i \wedge C_i)_{i=1}^n$ and a set of censoring indicators $(\delta_i)_{i=1}^n = (\mathbb{1}_{\{T_i \wedge C_i\}})_{i=1}^n$.

E Mini-batch sizing

The mini-batch sizing question is essential since it is a natural trade-off between computing time and precision. We know that computing $\nabla f_i(\theta)$ needs the computation of $|R_i| \in \{1, \ldots, n_{\text{pat}}\}$ inner products. Using the trick from Appendix B, one proves easily that computing a mini-batch $(1/n_{\text{mb}})\nabla f_{\mathcal{B}}(\theta)$ - where $\mathcal{B}$ is the set of n_{mb} index randomly picked - only needs $\max_{i \in \mathcal{B}} |R_i|$ inner products. A simple probability exercise gives us a key insight about the mini-batch size.
Let's assume that censoring is *uniform* over the set $\{1, 2, \ldots, n_{\text{pat}}\}$ meaning that $|R_i| = ci$ with $c > 1$. Then, we denote $u_1, u_2, \ldots, u_{n_{\text{mb}}} \sim \mathcal{U}[n]$ the indices independently sampled to compute the mini-batch i.e. $\mathcal{B} = \{u_i\}_{i=1}^{n_{\text{mb}}}$. Now we study the c.d.f. of $\max_{1 \leq i \leq n_{\text{mb}}} u_i$: for

$k \in \{1, 2, \ldots, n\}$,

$$\mathbb{P}\left(\max_{1\leq i\leq n_{\text{mb}}} u_i \leq k\right) = \prod_{i=1}^{n_{\text{mb}}} \mathbb{P}(u_i \leq k) = \left(\frac{\lfloor k \rfloor}{n}\right)^{n_{\text{mb}}},$$

$$\mathbb{P}\left(\max_{i\in\mathcal{B}} |R_i| \leq ck\right) = \left(\frac{\lfloor k \rfloor}{n}\right)^{n_{\text{mb}}},$$

$$\mathbb{P}\left(\max_{i\in\mathcal{B}} |R_i| \geq a\right) = 1 - \left(\frac{\lfloor a/c \rfloor}{n}\right)^{n_{\text{mb}}}, \text{ for } a < n_{\text{pat}}$$

The third equation leads us to consider $1 \ll n_{\text{mb}} \ll n$ to prevent both $\max_{i\in\mathcal{B}} |R_i|$ and $|\mathcal{B}|$ from being too large. This is why we used $n_{\text{mb}} = 0.1n$ or $n_{\text{mb}} = 0.01n$, depending of the size n of the dataset.